

# Challenges in Autonomous System Development

J. Connelly  
Institute for Defense Analyses,  
4850 Mark Center Drive,  
Alexandria, VA, USA  
jconnell@ida.org

W.S. Hong  
Institute for Defense Analyses,  
4850 Mark Center Drive,  
Alexandria, VA, USA  
whong@ida.org

R.B. Mahoney, Jr./D.A. Sparrow  
Institute for Defense Analyses,  
4850 Mark Center Drive,  
Alexandria, VA, USA  
rmahoney@ida.org  
dsparrow@ida.org

**Abstract**—The field of autonomous vehicles sits at the intersection of artificial intelligence (AI) and robotics, combining decision-making with real-time control. Autonomous vehicles are desired for use in search and rescue, urban reconnaissance, mine detonation, supply convoys, and more. The general adage is to use robots for anything dull, dirty, dangerous or dumb. While a great deal of research has been done on autonomous systems, there are only a handful of fielded examples incorporating machine autonomy beyond the level of teleoperation, especially in outdoor/complex environments. In an attempt to assess and understand the current state of the art in autonomous vehicle development, a few areas where unsolved problems remain became clear. This paper outlines those areas and provides suggestions for the focus of science and technology research. The first step in evaluating the current state of autonomous vehicle development was to develop a definition of autonomy. A number of autonomy level classification systems were reviewed. The resulting working definitions and classification schemes used by the authors are summarized in the opening sections of the paper. The remainder of the report discusses current approaches and challenges in decision-making and real-time control for autonomous vehicles. Suggested research focus areas for near-, mid-, and long-term development are also presented.

## I. INTRODUCTION

### A. Definition of Autonomy

What is autonomy? According to Webster [1], it is “the quality or state of being self-governing”. However, in the field of autonomous vehicles and military applications, autonomy is usually thought of as something more synonymous with “independence” or “intelligence”.

The official Department of Defense (DoD) definition of “autonomous operation”, from the DoD Dictionary of Military Terms, provides an interesting perspective on the concept and separates it somewhat from just autonomous vehicles:

“In air defense, the mode of operation assumed by a unit after it has lost all communications with higher echelons. The unit commander assumes full responsibility for control of weapons and engagement of hostile targets.” [2]

This definition also highlights the fact that autonomy does not apply only to machines, but is already a working concept within the military chain of command. Therefore, when considering autonomy, the terms “Authority” and “Agent” instead of “human” and “computer” are suggested. In this way, the discussions are not limited to the hierarchy as it is

currently envisioned.

One interesting characterization of autonomy found was “[autonomy] is whatever we don’t know how to do yet. Once we know how to do it, we call it an algorithm.”<sup>1</sup> In fact, this is more widespread today than generally realized. Some functions taken for granted in cars or planes today make and execute decisions independently and thus may be considered autonomous subsystems, e.g. optimization of fuel and battery power consumption ratios in hybrid vehicles, air bags, and anti-lock brakes. However, because the whole car is not autonomous, there is a tendency to minimize the successes that have been attained thus far, and characterize them as “automatic” rather than “autonomous”.

What is the difference between “automatic” and “autonomous”? One distinction may be to say that something automatic has only one “choice” between two possible states, e.g. ‘on’ or ‘off’. Another classification would say that automatic systems take in only one input for making the decision. In either case, current air bags and anti-lock brakes would likely fall in the “automatic” instead of “autonomous” category. Autonomous systems could then be ones that process multiple inputs before acting, e.g. a braking system that considers both wheel slippage and speedometer measurements and only deploys if the car is traveling faster than 30mph. Alternatively, autonomous systems may be those that have more than two possible states, and so have to make more than an “on/off” choice.

The relative merits of differing distinctions between “autonomous” and “automatic” are hard to measure—there are continuing debates and the presence of counterexamples in any classification system or definition proposed to date. If the line between “automatic” and “autonomous” is drawn based on number of choices or whether the system is following rules instead of “making its own decisions”, then any current system would be considered “automatic”, not “autonomous”, because they are all deterministic in their decision making. This observation raises the question of whether any currently foreseeable (i.e. deterministic) system is truly autonomous; the ambiguity of the term may be why many sectors are choosing to use the term “unmanned” instead. However, in order to encourage research and development in useful

---

<sup>1</sup> Patrick Winston, former director of MIT’s Artificial Intelligence Laboratory, as quoted in “Autonomous Land Vehicles” by Dr. Hugh Durrant-Whyte.

near-term areas, one might use “autonomy” in an inclusive rather than restrictive sense. Therefore, we propose the following working definition of an autonomous system:

*An autonomous system is one that makes and executes a decision to achieve a goal without full, direct human control.*

Here “system” does not have to mean an entire vehicle; it could also mean a subsystem like the anti-lock brake system (ABS) example. By this definition, automatic is not distinct from autonomous, but is a subset instead. This inclusive definition dovetails nicely with the ongoing efforts to classify “levels of autonomy”. These levels would depend on such things as mission complexity or level of required human interaction. Automatic systems (single input to single output) would occupy the lower end of any autonomy scale.

In developing this working definition of autonomy, it became clear that there are two main areas of development for an autonomous vehicle: decision-making and real-time control. Generally speaking, the decision-making side corresponds to “autonomous” (or independence) and the real-time control corresponds to “vehicle” (or execution), although the line between the two can be a bit fuzzy at times. There is clearly some local decision-making that takes place within the realm of real-time control, such as in local navigation and obstacle avoidance. Otherwise, robots would run into an obstacle while trying to decide whether to go left or right around it. Similarly, these two categories cannot stand alone. Developers of autonomous vehicles cannot work on autonomy and computer processing separately from working on vehicle mechanics—the integration of these two areas into one physical system presents a significant challenge in and of itself. Not only does the computer equipment need to be able to physically withstand the operational environment onboard a moving vehicle, but it also needs to appropriately connect the algorithms to the incoming sensor data and decide which sensor information is needed in the first place.

Consistent with all three definitions above, “autonomy” is not a technology itself, but rather a capability enabled by supporting technologies. Dr. Durrant-Whyte divides these technologies into five categories: mobility, localization, navigation, planning, and communication [3]. Mobility includes the real-time control and mechanics of the vehicle itself. Localization incorporates sensors and software to identify the vehicle’s position, attitude, velocity, and acceleration. Navigation, to include local obstacle avoidance, combines decision-making and real-time control. Planning includes mission- and task-level decisions, waypoint generation, task allocation, etc. Communication involves all the links between the vehicle and teammates, operators, and command and control. These five categories summarize the main contributing technology areas for autonomous vehicles.

### B. Autonomy Levels

There has been extensive work by others attempting to go beyond a working definition of autonomy to quantified autonomy levels [5–8]. From these various scales, four main categories were identified: piloted vehicle, authority in the

loop, authority on the loop, and authority out of the loop. These categories are based on work presented by Chad Hawthorne and Dave Scheidt at the Johns Hopkins University Applied Physics Lab [9]. These categories will differ qualitatively in engineering approach, test and evaluation activities, and demonstrated reliability prior to fielding.

## II. DECISION-MAKING

Autonomous decision-making is an incredibly complex subject, especially given the fact that scientists do not fully understand how the human brain works and makes decisions. There appear to be two main categories of machine decision-making: reduction and learning. Figure 1 below reveals just how complex autonomous decision-making processes can be [10]. Indeed, one way to measure levels of autonomy would be to consider how many layers of the decision-making process portrayed are employed by the unmanned system.

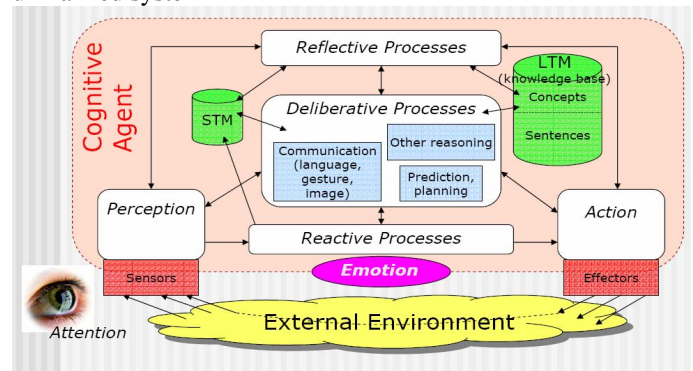


Fig. 1: Diagram of a Cognitive Agent

### A. Reduction

As illustrated above, decision-making can involve much more than a simple binary selection. Humans incorporate a priori knowledge, context, and emotions when making decisions. In the reduction approach to autonomous decision-making, those elements are largely excluded. Instead, the problem is reduced to a simple, clearly-defined input-output mapping.

Automotive subsystems provide numerous examples of this approach. Anti-lock brake systems have long been the standard in American cars. Anti-lock brakes use a sensor that detects changes in wheel spin rate. When that sensor readout passes a certain threshold, the automatic brake is activated. There is a direct mapping of input to output, a clear rule for which action to take and only two choices for action: activate or not. ABS incorporate both the autonomous decision-making mentioned above and real-time control, in the pumping of the brake. From the definition of an autonomous system proposed above, the “goal” declared by the human driver is to stop; the ABS then decides how to accomplish that goal—whether the pumping is required in the situation—and then executes that decision, all more quickly than a human driver could.

A similar threshold sensor with a binary output option found in automobiles is the air bag. When deceleration is faster than a certain limit, the air bag deploys. Again, the air bag is an autonomous subsystem—although in this case, even more control is ceded to the computer, because the driver cannot override the decision to deploy just by lifting his foot up the way he can with his brakes.

An extension of this method is used in the new hybrid cars to optimize the ratio between fuel and battery power consumption as a function of speed, remaining battery life, etc. In this situation there are more dimensions than for ABS or airbags: the onboard computer needs to determine the optimal split between combustion and battery power and to execute the switching back and forth.

An example of the reduction approach that has already been applied in robotics is simple obstacle avoidance. The problem can be reduced to a binary output—“can I go straight or not?” There are only two output options and potentially only one required sensor. This is a very simplified method and would probably not detect things like cliffs or chain link fences, depending on the capabilities and sensitivity of the vehicle’s sensors, but it can be enough to successfully avoid obstacles in indoor or relatively uniform outdoor environments.

A variation on the reduction approach is the use of multi-robot systems. The concept is to give simple tasks/capabilities to each robot and connect them via a wireless network. By separating the overall mission into smaller subtasks, the complexity of the problem has been reduced to one that can be physically accomplished by current robots.

If the problem is more complex and a simple mapping is not obvious, researchers can conduct experiments, collect data, and write algorithms that characterize the domain within a given area (e.g. the flight envelope for an airplane autopilot). Then computing power can be employed to perform the bookkeeping and keep track of sensor data, the aerodynamic effects on the vehicle, etc.

A similar “bookkeeping” approach has been suggested for obstacle avoidance. Ideally, this method would help optimize a ground robot’s route—different maximum safe speeds could be connected with each terrain type, for example. However, building the database would be quite tedious and require some foreknowledge of the route. This approach degrades in changing terrain, and may require storage and processing beyond space/weight/power limits.

One issue with the reduction approach is that the “rules” given to the computer are only good within the given operational envelope—it is very difficult to cope with scenarios that fall outside the bounds of predicted patterns. For example, in 2001 a P-3 was involved in a mid-air collision with a Chinese aircraft and the pilot managed to land the plane safely [11]. To accomplish this feat, the pilot had to assess the plane’s changed response with enough speed and accuracy to prevent the plane from crashing. Current autopilots, such as that on the Global Hawk that recently landed safely after an engine flameout [12], may be able to

recover from types of in-flight failure that have standard responses that can be programmed in ahead of time. However, other types of failure may be too far outside the operational capability of the aircraft, requiring human-level experience, intuition, and rapid learning in order to successfully recover.

### *B. Learning*

The other approach is to attack complex, incompletely characterized problems with superior computing power. The example of the P-3 pilot recovering from a midair collision is exactly the type of learning that the AI world is trying to recreate in order to tackle complex, incompletely-characterized problems. Missions beyond a certain level of complexity may never be possible without some leap in computer learning. For example, with the terrain database approach discussed above, it seems implausible to develop a database with any significant operational envelope for an uncertain or unknown outdoor environment. When the environment is structured or can be structured without disruption, it becomes possible to more fully characterize it, and achieve mission success with modest machine learning. Robots in manufacturing plants that follow lines or magnets in the floor are an example of such an application.

Robots need structure; that is how the variation and surprise can be restricted to levels that current processing power and algorithms can handle. Vehicle developers often find a way to bring structure to the environment and make it navigable for the unmanned systems. However, many of the environments in which users would like to send robots, such as natural disasters or military operations, cannot be structured ahead of time. If managing the environment is infeasible, it becomes increasingly important to develop learning capabilities so that robots can function in changing or unknown environments. The Defense Advanced Research Projects Agency (DARPA) has a number of programs focused on advancing machine learning and autonomous decision-making, such as Learning Applied to Ground Robots (LAGR) and Real World Reasoning (REAL). However, these programs are still in the early research phases and lie outside the scope of this report. To put the challenges these programs face in context, a first responder or soldier has 18 years or more of learning, before the task specific training starts.

### *C. Summary*

Mission complexity for fully autonomous systems will be severely limited until significant AI developments are achieved. However, there are still a number of useful steps that could be taken, and it is in these areas that research and development would be most useful in the near-term. High payoff pursuits for near-term development include:

- Further characterizing the environment, i.e. quantifying and expanding the understood operational envelope for ground vehicles
- Increasing reliability of communication links in order to progress from tethered teleoperation to wireless
- Making sensible choices about the role and application

of autonomous vehicles and focusing development on those applications,

- Building machines robust enough to withstand less fine-toothed decision-making

### III. REAL-TIME CONTROL

Real-time control concerns, in part, the physical aspects of an autonomous vehicle, as well as the translation from decision to action. Decision-making is still largely regarded as “science” and the real-time control is primarily considered “engineering”. However, this does not mean that all the unsolved problems are on the decision-making side and that successful real-time control is just a matter of working out some engineering details.

One continually difficult problem is local navigation and obstacle avoidance. Vehicles need to fuse and process sensor data at fast enough speeds and with enough accuracy to prevent running into things or getting stuck before higher-level decisions can be made. In a way, obstacle avoidance captures both real-time control and decision-making, albeit on the small-scale, local level. Current appropriate sensor packages are few and far between. While the problem may be “solved” in a performance sense, if the sensor that has been developed does not meet space, weight, power, and cost constraints, then that sensor is not a solution at all. Because the work done in this field is so application-specific, there appear to be numerous individual claims of solutions or successful demonstrations. Yet those successes do not readily translate to other programs or platforms. Therefore, it would be premature to consider such issues “solved” problems.

Much of the difficulty in developing autonomous vehicles capable of complex missions is that researchers don’t understand how humans make decisions or perform those same tasks. The same is true for some aspects of real-time control. The human hand is an incredibly complex array of sensors and interconnected effectors. The sensitivity of force sensors in our fingers is unparalleled. There is also a certain amount of local processing that takes place—for example, if a person touches a hot stove, his hand jerks away before the brain has even had time to register that the surface was hot. Similarly, if someone walks into a door frame, they don’t break a shoulder; they automatically start reducing the pressure applied at the point of contact. A robot, on the other hand, can snap an appendage off if it runs into a doorframe or tries to find a light switch and flip it on in a dark room. So there is a tradeoff between sensitivity and precision. The current sensor packages available for autonomous vehicles provide much less information to the decision-making algorithm than humans use on a regular basis. While building a humanoid robot may not be a primary interest for the military, this example highlights one of the significant limiting factors in the application of robots. Therefore, the best focus for development efforts is on tasks at which robots exceed human performance, rather than ones that just try to mimic humans.

A final challenge facing autonomous vehicle development from the real-time control side is systems integration. It is essential that all the components be mounted on board a mission-appropriate vehicle and that they survive the mission. Current sensor packages are generally too expensive or too bulky for practical applications—especially on ground vehicles. The vehicle also needs to be robust enough to protect all of the sensor and computing equipment when navigating in rough terrain. Similarly, a highly advanced sensor may be developed that would allow for significantly increased autonomy, but if that sensor requires a massive power supply, the vehicle would not be able to move very far from the base station. The systems integration challenge highlights a key issue in future autonomous vehicle development—specialization vs. generalization. While general programs and packages applicable across platforms appear to be the ideal, truly successful robots to date have been developed for specific missions. The specialized approach limits the systems integration issues, because the pieces are designed to go together more readily. While a common architecture or sensor platform may be on the research horizon, for the near-term, the field might be better served to focus development on more capable, task-specific vehicles.

### IV. CONCLUSION

Research efforts in AI and cognitive computing have been largely theoretical or simulation-based. There is a disconnect between the field of robotics and the field of cognitive computing, especially when it comes to real-world implementation. Current artificial intelligence research is, by and large, not being designed for implementation on board a moving vehicle; yet robots will only be able to achieve a certain minimal level of complexity without integrating AI concepts and developments. If any significant advances beyond teleoperation are to be made in autonomous vehicle development, these two research fields need to come together and use advances from each area in the development of new vehicles.

Up to this point, autonomous vehicle development has been either highly application specific or too theoretical to apply on board an actual vehicle. There is a commonly-held hope that a single architecture or navigation method could be developed that would apply across platforms or applications, but that does not appear to be an option in the near term. Basic research should continue to provide new capabilities. However, it seems that there are many factors specific to each mission and/or environment that require specific development efforts for both the decision-making approach and the real-time control for each application. Thus far, the more useful a vehicle has been, the more specialized its development was.

Programs that focus on real-world implementation appear to be having more success. Their progress along the autonomy scale may be modest, but they are fielded and saving the lives of disaster victims or soldiers. However, DoD acquires

general-purpose equipment precisely because it is more difficult to anticipate operational needs than commercial ones. Therefore, the primary goal for the Department of Defense seems to be to increase the mission complexity and environmental variability in which unmanned vehicles are capable of performing. In this way human soldiers can be removed from dangerous, dull, dirty, and dumb situations. For DoD applications, at least, this motivates advances in cross platform commonality and in autonomy.

Basic AI research is still required, especially in the area of transfer learning—generalizing from a previous example to a novel situation. Until this trait of humans is more fully understood and accomplished in computers—or its effects mimicked—there will continue to be long training times and high costs, often resulting in brittle performance. In the slightly closer-term, integrating AI systems on board robotic platforms would yield enormous payoffs. We have only begun to focus on what happens with AI systems in real world environments. Incorporating context and intuition into machine systems—or at least modeling and understand their role in human decision-making processes well enough to assess the impact of their absence in autonomous vehicles warrants additional attention is both AI and psychology .

Outdoor obstacle avoidance remains a key issue for ground vehicles and is probably the area that already incorporates significant AI but also runs into the most problems due to the incredible variability of the terrain. Obstacle avoidance is much more straightforward for an unmanned aerial vehicle (UAV): there are far fewer obstacles above tree level and less variation in the environment. Indoor environments and highly structured outdoor environments such as those in agricultural applications are more tractable, specifically because structure has been imposed on the environment.

#### REFERENCES

- [1] "Autonomy." *Merriam-Webster's Collegiate Dictionary*, 9th ed, 1986
- [2] United States. Department of Defense. Defense Technical Information Center. DoD Dictionary of Military Terms. Jul. 2005 <<http://www.dtic.mil/doctrine/jel/doddict/data/a/00599.html>>
- [3] Durrant-Whyte, Hugh. "Autonomous Land Vehicles", Proc. IMechE. Vol. 219 Part I: J. Systems and Control Engineering. IMechE 2005
- [4] Spence, Floyd D., National Defense Authorization Act for Fiscal Year 2001 (P.L. 106-398 Sec. 220)
- [5] Clough, Bruce T. "Metrics, Schmetrics! How The Heck Do You Determine a UAV's Autonomy Anyway?", 2002. Jun. 2005 <[http://www.isd.mel.nist.gov/research\\_areas/research\\_engineering/Performance\\_Metrics/PerMIS\\_2002\\_Proceedings/Clough.pdf](http://www.isd.mel.nist.gov/research_areas/research_engineering/Performance_Metrics/PerMIS_2002_Proceedings/Clough.pdf)>
- [6] United States National Institute of Standards and Technology, "ALFUS Framework", 2005. Jun. 2005 <[http://www.isd.mel.nist.gov/projects/autonomy\\_levels/](http://www.isd.mel.nist.gov/projects/autonomy_levels/)>
- [7] Sheridan, Thomas B. and Verplank, William L., "Human and Computer Control of Undersea Teleoperations", 1978
- [8] IDA Rosetta Stone Paper
- [9] Hawthorne, Chad and Scheidt, Dave, "Moving Emergent Behavior Algorithms from Simulation to Hardware: Command and Control of Autonomous UxV's", 10th International Command and Control Research and Technology Symposium, 2005
- [10] U.S. Department of Defense, Defense Advanced Research Projects Agency, <<http://www.darpa.mil/ipto/briefings/IPTO-Overview.pdf>>
- [11] Venik's Aviation, "Midair Collision Over China", 2001, 10 JAN 2006,

<<http://www.aeronautics.ru/news/news001/news031.htm>>

[12] Harvey, David S. "Global Hawk: Flameout Led To Automatic Afghan Alternate", SEP 2005, 10 JAN 2006 <<http://uvscanada.org/blog/?p=46>>